arXiv:1411.7920v2 [math.PR] 4 Dec 2014

Submitted to the Annals of Probability

PROBABILITY THEORY WITHOUT BAYES' RULE

By Samuel G. Rodriques[*,†]

Department of Physics, Massachusetts Institute of Technology

Within the Kolmogorov theory of probability, Bayes' rule allows one to perform statistical inference by relating conditional probabilities to unconditional probabilities. As we show here, however, there is a continuous set of alternative inference rules that yield the same results, and that may have computational or practical advantages for certain problems. We formulate generalized axioms for probability theory, according to which the reverse conditional probability distribution $P(B|A)$ is not specified by the forward conditional probability distribution $P(A|B)$ and the marginals $P(A)$ and $P(B)$. Thus, in order to perform statistical inference, one must specify an additional "inference axiom," which relates $P(B|A)$ to $P(A|B)$, $P(A)$, and $P(B)$. We show that when Bayes' rule is chosen as the inference axiom, the axioms are equivalent to the classical Kolmogorov axioms. We then derive consistency conditions on the inference axiom, and thereby characterize the set of all possible rules for inference. The set of "first-order" inference axioms, defined as the set of axioms in which $P(B|A)$ depends on the first power of $P(A|B)$, is found to be a 1-simplex, with Bayes' rule at one of the extreme points. The other extreme point, the "inversion rule," is studied in depth.

1. Introduction. Let A and B be random variables taking values $a_1, \ldots, a_{N_A}$ and $b_1, \ldots, b_{N_B}$, respectively, and let $P(a_i, b_j)$ be the joint probability distribution for A and B, which may be represented as a matrix $\mathbf{P}(A, B)$ of size $N_A \times N_B$. The conditional probability distribution $P(a_i|b_j)$ may likewise be represented as an $N_A \times N_B$ matrix, denoted $\mathbf{P}(A|B)$, with the i,jth component given according to the standard Kolmogorov axioms of probability Kolmogorov [1933] by

$$P(a_i, b_j) = P(a_i|b_j)P(b_j). \tag{1}$$

In the Kolmogorov framework, $P(a_i, b_j)$ is viewed as the measure of the intersection of sets representing a_i and b_j. It thus follows from the symmetry of the intersection operator that $P(a_i, b_j) = P(b_j, a_i)$. From the definition

[*]The author thanks Maxim Rabinovich, Sina Tootoonian, David Turban, and Steven Lindell for helpful discussions and feedback.

[†]The author gratefully acknowledges the support of the Winston Churchill Foundation of the United States, and the Fannie and John Hertz Foundation.

of conditional probabilities we can thus derive Bayes' rule, which gives the relationship between $P(A|B)$ and $P(B|A)$:

$$P(a_i|b_j) = \frac{P(b_j|a_i)P(a_i)}{P(b_j)}. \tag{2}$$

However, Bayes' rule gives rise to some surprising behavior. By unitarity and additivity, it follows that

$$P(a_i) = \sum_j P(a_i, b_j), \tag{3}$$

and hence

$$P(a_i) = \sum_j P(a_i|b_j)P(b_j). \tag{4}$$

Equation 4 is simply the statement that the vector $\vec{P}(A)$ with components $P(a_i)$ is obtained from the vector $\vec{P}(B)$ with components $P(b_j)$ by the action of $\mathbf{P}(A|B)$. Likewise for B, we have

$$P(b_j) = \sum_i P(b_j|a_i)P(a_i). \tag{5}$$

By substitution, it follows that

$$P(b_k) = \sum_{ij} P(b_k|a_i)P(a_i|b_j)P(b_j). \tag{6}$$

Given $\mathbf{P}(A|B)$, equation 6 must hold for all choices of the probability distribution $\vec{P}(B)$. One natural conclusion would be that $\mathbf{P}(A|B)$ and $\mathbf{P}(B|A)$ are inverses, i.e.,

$$\sum_i P(b_k|a_i)P(a_i|b_j) = \delta_{kj}. \tag{7}$$

However, the matrix inverse of $\mathbf{P}(A|B)$ will have negative values and values greater than 1 in general. Instead, given constant $\mathbf{P}(A|B)$, $\mathbf{P}(B|A)$ varies with $\vec{P}(B)$ according to Bayes' rule, in such a way that $\vec{P}(B)$ is always an eigenvalue of the matrix $\mathbf{P}(B|A)\mathbf{P}(A|B)$ with eigenvalue 1. This situation is particularly dissatisfying when one thinks of $\mathbf{P}(A|B)$ as describing the action of a channel, in which case $\mathbf{P}(B|A)$ would describe the result of running the channel backwards. In this case, Bayes' rule does not allow a simultaneous description of the forward channel and the reverse channel independent of

the distribution that goes through it. One imagines that the matrix inverse might be a more appropriate description of the reverse channel than the matrix $\mathbf{P}(B|A)$ obtained via Bayes' rule.

The matrix inverse of $\mathbf{P}(A|B)$ shows up in other ways as well. For example, suppose we are given $\vec{P}(A)$ and $\mathbf{P}(A|B)$, and wish to infer $\vec{P}(B)$. This task is not possible using Bayes' rule, since $\mathbf{P}(B|A)$ depends on $\vec{P}(B)$; one requires a prior distribution. Nonetheless, $\vec{P}(B)$ may always be obtained from $\vec{P}(A)$ via the inverse of the matrix $\mathbf{P}(A|B)$:

$$\vec{P}(B) = \mathbf{P}(A|B)^{-1}\vec{P}(A). \tag{8}$$

Within the strict framework of Kolmogorov probability theory, one must view this procedure as a coincidence, since the matrix $P(A|B)^{-1}$ is not a conditional probability distribution.

As we show here, the application of $\mathbf{P}(A|B)^{-1}$ in this way may be explained using an alternative axiomatization of probability theory in which Bayes' rule is replaced by the requirement that

$$\mathbf{P}(B|A) = \mathbf{P}(A|B)^{-1}, \tag{9}$$

referred to as the "inversion rule." Moreover, the inversion rule is only one of a continuous set of possible alternative axiomatizations of probability theory, each of which comes with its own rule for statistical inference. In the second section, we provide the general axiomatic framework for probability theory that we will consider in the remainder of the paper. In the third section, we explore the specific axiom that gives rise to equation (9) as a rule for inference, and show how it may be used alongside Bayes' rule in practice. Finally, in the fourth section, we characterize the set of all possible inference rules. We find that Bayes' rule and the inversion rule form a conjugate pair, and that they can be used to generate all other rules for inference.

2. Axioms. We seek to formulate a set of axioms that preserves the notion of a probability distribution in its entirety without Bayes' rule. Since Bayes' rule is a consequence of the symmetry of the intersection operator, it is necessary to replace Kolmogorov's axiomatization in terms of set theory by a similar axiomatization in terms of sequences. The primary hurdle will be to establish analogs of the union and intersection operators for sequences, which is necessary to facilitate comparison with the Kolmogorov axioms.

We consider the case of N random variables $A^{(1)}, \ldots, A^{(N)}$, where $A^{(i)}$ takes values $a^i_j \in J_i$, for $j = 1, \ldots, M_i$. Let S be the set of all permutations of the integers from 1 to N. Given some $s \in S$, we denote by E_s the set of all sequences Q of N elements such that the ith element in Q is in $J_{s(i)}$. Q

is then said to have order s. Given sequences Q and Q', Q is said to be a subsequence of Q' if Q may be obtained by removing elements from Q'. We denote by F_s the set of all subsequences of elements of E_s. A sequence Q is then said to have order s if it is in F_s. For any sequence Q, we denote by $\bar{Q}$ the set of elements appearing at some position in Q, referred to as the "membership set" of Q. We write $Q \subseteq_s Q'$ if Q and Q' both have order s, and if $\bar{Q} \subseteq \bar{Q}'$.

For any $Q \in F_s$, we define $R_s(Q)$ to be the set of sequences in E_s of which Q is a subsequence, i.e.,

$$R_s(Q) \equiv \{Q' \in E_s | Q \subseteq_s Q'\}. \tag{10}$$

We may now evidently define the operator $\cap_s$ such that for all $Q_1, Q_2 \in F_s$,

$$Q_1 \cap_s Q_2 = \{Q' \in R_s(Q_1) \cap R_s(Q_2)\}. \tag{11}$$

Likewise, we define the operator $\cup_s$ such that for all Q_1, $Q_2 \in F_s$,

$$Q_1 \cup_s Q_2 = \{Q' \in R_s(Q_1) \cup R_s(Q_2)\}. \tag{12}$$

The axioms are then as follows, and are formulated for the greatest similarity to the original axioms of Kolmogorov:

Axiom 2.1. For all $s \in S$ and for all $Q \in E_s$, $P(Q)$ is a real-valued function referred to as the probability of Q.

Axiom 2.2. For all $s \in S$ and for all $Q \in F_s$,

$$P(Q) = \sum_{Q' \in R_s(Q)} P(Q'). \tag{13}$$

Crucially, we note that because P is not explicitly a function of s, axiom 2.2 implies that for all s,$s' \in S$ and for all $Q \in F_s \cap F_{s'}$,

$$\sum_{Q' \in R_s(Q)} P(Q') = P(Q) = \sum_{Q' \in R_{s'}(Q)} P(Q'). \tag{14}$$

Equation 14 is equivalent to the statement that the marginal probability distribution over a subset of variables depends only on the marginal ordering of the variables in the subset.

Axiom 2.3. For all $s \in S$, we have

$$P(E_s) = 1. \tag{15}$$

AXIOM 2.4. For all $Q_1, Q_2 \in F_s$ such that $R_s(Q_1) \cap R_s(Q_2) = \phi$,

$$P(Q_1 \cup_s Q_2) = P(Q_1) + P(Q_2). \tag{16}$$

For any given value of s, the axioms 2.1 through 2.4 are identical to the standard statement of Kolmogorov's axioms without the requirement that $P(Q)$ be non-negative. Thus, the axioms are consistent, but not complete.

For the sake of concreteness, we will henceforth focus on the case of two variables, A and B, so that there are only two orderings. Then a conditional probability may be easily defined[1]:

$$P(a_i|b_j) \equiv \frac{P(a_i, b_j)}{P(b_j)}. \tag{17}$$

Note that the ordering of the variables should be read off from right to left, so $P(a_i, b_j)$ is the probability of observing a_i and b_j relative to the ordering in which B precedes A. This convention has been chosen to keep with the standard convention for conditional probabilities, according to which the conditioning variables are written on the right, and the conditioned variables are written on the left.

We are interested in our ability to perform statistical inference within this system of axioms. In the problem of statistical inference, we consider an observable variable A and a hidden variable B. The goal is to determine a posterior for the hidden variable, $P(b_j|a_i)$, in terms of a model, $P(a_i|b_j)$, and a prior, $P(b_j)$. Within the current axiomatic system, the posterior distribution $P(b_j|a_i)$ is constrained in terms of $P(a_i|b_j)$ and $P(b_j)$ via equation 14, which applied to the current system becomes

$$\sum_j P(a_i, b_j) = P(a_i) = \sum_j P(b_j, a_i). \tag{18}$$

However, this constraint does not specify $P(b_j|a_i)$ uniquely in terms of $P(a_i|b_j)$. Thus, in order to perform statistical inference, it is necessary to impose an additional "inference axiom" that provides us with a rule for inferring $P(b_j|a_i)$ in terms of $P(a_i|b_j)$ and $P(b_j)$. Bayes' rule is one possible rule of inference, and may be recovered from the current axiomatization by requiring that the joint probability over some set of variables is independent of the ordering of those variables[2], in which case $P(a_i, b_j) = P(b_j, a_i)$ for all

[1] The conditional probability may be easily defined in the same way for any number of variables, but the notation required to treat the general case is very cumbersome, so has been avoided here.

[2] In the general formalism, the axiom is that for all $Q \in E_s$ and $Q' \in E_{s'}$ such that $\bar{Q} = \bar{Q}'$, $P(Q) = P(Q')$.

i and j, and hence

$$P(b_j, a_i) = P(a_i|b_j)\frac{P(b_j)}{P(a_i)}. \tag{19}$$

As we show below, however, there are many other choices of inference axioms that will nonetheless yield correct results.

In all that follows, we will assume that $\mathbf{P}(A|B)$ is invertible, which ensures that the problem of inference will be uniquely solvable.

3. The Inversion Rule. Before examining the set of all possible inference axioms, it will be useful to examine one particular inference rule, the inversion rule, in depth. The inversion rule is given by equation (9), and is noteworthy for being the only inference rule in which $\mathbf{P}(B|A)$ is specified exclusively in terms of $\mathbf{P}(A|B)$. It is thus associated with the following inference axiom:

AXIOM 3.1. $\mathbf{P}(A|B)$ and $\mathbf{P}(B|A)$ may be specified independently of $\vec{P}(B)$ and $\vec{P}(A)$.

We prove that this axiom uniquely specifies the inversion rule as follows:

THEOREM 1 (Inversion Theorem). *For all non-product distributions $P(A, B)$, the matrices $\mathbf{P}(A|B)\mathbf{P}(B|A)$ and $\mathbf{P}(B|A)\mathbf{P}(A|B)$ are equal to the identities of size N_A and N_B respectively.*

PROOF. From the definition of conditional probability, we have

$$P(a_i) = \sum_j P(a_i|b_j)P(b_j), \tag{20}$$

and

$$P(b_j) = \sum_k P(b_j|a_k)P(a_k). \tag{21}$$

From eq. (14), we have

$$P(a_i) = \sum_{j,k} P(a_i|b_j)P(b_j|a_k)P(a_k). \tag{22}$$

Eq. (22) may be recast in matrix form as

$$\vec{P}(A) = \mathbf{P}(A|B)\mathbf{P}(B|A)\vec{P}(A). \tag{23}$$

Equation 23 must hold for all choices of $\vec{P}(A)$. However, $\mathbf{P}(A|B)\mathbf{P}(B|A)$ is independent of $\vec{P}(A)$ by axiom 3.1. It follows that $\mathbf{P}(A|B)\mathbf{P}(B|A)$ must be the identity. The proof for $\mathbf{P}(B|A)\mathbf{P}(A|B)$ is identical. □

Equation (23) may in general still be applied even if $\mathbf{P}(A, B)$ does not have full rank by restricting the inversion to the support of $\mathbf{P}(A, B)$. However, when $\mathbf{P}(A, B)$ is a product distribution, all quantities associated with all other orderings of the variables are undefined. In the case of a product distribution, it is impossible to decouple $\mathbf{P}(A|B)$ from $\vec{P}(A)$ since $P(a_i|b_j) = P(a_i)$ for all j, so axiom 3.1 is inconsistent with the other axioms.

The joint distributions for the two different orderings may be related as follows:

THEOREM 2. *Let $\mathbf{P}(A, B)$ be a non-product joint distribution. Then,*

$$\mathbf{P}(B, A) = \mathbf{P}(B)\mathbf{P}(A, B)^{-1}\mathbf{P}(A). \tag{24}$$

PROOF. Combining (23) with the definition of conditional probability, it is clear that

$$\mathbf{P}(A|B) = \mathbf{P}(A, B)\mathbf{P}(B)^{-1} \tag{25}$$

and

$$\mathbf{P}(B|A) = \mathbf{P}(B, A)\mathbf{P}(A)^{-1}. \tag{26}$$

Inverting (25) and equating the left-hand sides of the two equations, we obtain

$$\mathbf{P}(B, A)\mathbf{P}(A)^{-1} = \mathbf{P}(B)\mathbf{P}(A, B)^{-1},$$

from which it follows that

$$\mathbf{P}(B, A) = \mathbf{P}(B)\mathbf{P}(A, B)^{-1}\mathbf{P}(A). \tag{27}$$

□

Clearly, if all the entries of $\mathbf{P}(A, B)$ are positive then the entries of $\mathbf{P}(B, A)$ will neither be strictly positive nor smaller than 1 in general.

A further interesting consequence of Theorem 2 is that the function P is completely specified by the joint distribution for only a single order, and hence contains the same amount of information as the joint distribution in the Kolmogorov framework. More generally, as long as the chosen inference axiom specifies $\mathbf{P}(B|A)$ completely in terms of $\mathbf{P}(A|B)$ and $\mathbf{P}(B)$, $\mathbf{P}(B, A)$ will be completely specified in terms of $\mathbf{P}(A, B)$. For this reason, we are always free to choose a preferred ordering of the variables, relative to which all joint, conditional, and marginal probabilities agree with the Kolmogorov probabilities:

COROLLARY 1. *Let P_K be a non-product probability distribution obeying the Kolmogorov axioms of probability. Then there exists a probability distribution P_I obeying axioms 2.1 through 2.4 and 3.1 such that for all i and j, $P_K(a_i, b_j) = P_I(a_i, b_j)$.*

Corollary 1 allows us to use the inversion rule to perform statistical inference on Kolmogorov probabilities as follows. Given a model $\mathbf{P}_K(A|B)$ and an empirical estimate $\tilde{P}_K(A)$ for $\vec{P}_K(A)$, we consider the probability distribution $\mathbf{P}_I(A, B)$ in the inversion framework, defined such that $\mathbf{P}_I(A, B) = \mathbf{P}_K(A, B)$. Then, the best empirical estimate for $\tilde{P}_I(B)$ using the inversion rule is given by

$$\begin{aligned} \tilde{P}_I(B) &= \mathbf{P}_I(B|A)\tilde{P}_I(A) \\ &= \mathbf{P}_I(A|B)^{-1}\tilde{P}_I(A) \\ &= \mathbf{P}_K(A|B)^{-1}\tilde{P}_K(A) \end{aligned} \tag{28}$$

Furthermore, in the limit as $\tilde{P}_K(A) \to \vec{P}_K(A)$, the estimate $\tilde{P}_I(B)$ obviously converges to $\vec{P}_I(B)$. Because $\vec{P}_I(B) = \vec{P}_K(B)$, we may interpret $\tilde{P}_I(B)$ as an approximation to $\tilde{P}_K(B)$:

$$\tilde{P}_K(B) = \mathbf{P}_K(A|B)^{-1}\tilde{P}_K(A). \tag{29}$$

Equation 29 cannot be justified on the basis of the axioms of Kolmogorov probability theory. In the Kolmogorov framework, equation 8 can only be used to relate the specific distributions $\vec{P}_K(A)$ and $\vec{P}_K(B)$. Because $\mathbf{P}_K(A|B)^{-1}$ is not a conditional probability distribution, there is no guarantee that $\tilde{P}_K(B)$ obtained using (29) is a probability distribution. Nonetheless, using the inversion framework, we have proven the validity of (29) as a method of performing inference.

One interpretation of the preceding discussion is that it is possible to use the inversion rule to infer a prior from the data. Surprisingly, because the inversion rule and Bayes' rule are independent, one may then apply Bayes' rule to do Bayesian inference using the inferred prior. We have

$$\begin{aligned} \tilde{\mathbf{P}}_K(B|A) &= \tilde{\mathbf{P}}(B)\mathbf{P}_K(A|B)^T\tilde{\mathbf{P}}(A)^{-1} \\ &= \mathrm{Diag}\left[\mathbf{P}_K(A|B)^{-1}\tilde{P}(A)\right]\mathbf{P}_K(A|B)^T\tilde{\mathbf{P}}(A)^{-1} \end{aligned} \tag{30}$$

Here, Diag $[\vec{v}]$ is the diagonal matrix with the entries of $\vec{v}$ along the diagonal. The estimate $\tilde{\mathbf{P}}_K(B|A)$ converges to $\mathbf{P}_K(B|A)$ as $\tilde{P}(A) \to \vec{P}(A)$, so $\tilde{P}(B)$ is guaranteed to be a "good" prior, even if it is not a Kolmogorov probability distribution.

4. All Possible Inference Axioms. We now examine the set of all possible inference axioms that can be added to axioms 2.1 through 2.4. An inference axiom is any way of specifying $\mathbf{P}(B|A)$ in terms of $\mathbf{P}(A|B)$ and $\vec{P}(B)$. For an inference axiom to be consistent with axioms 2.1 through 2.4, the matrix $\mathbf{P}(B|A)$ must satisfy equation 18. For equation 18 to be satisfied, the columns of $\mathbf{P}(B|A)$ must sum to 1, and the marginal distributions must not depend on the order of the joint distribution over which one chooses to marginalize. From this latter requirement and the definition of conditional probabilities, we derive the requirement that

$$\vec{P}(B) = \mathbf{P}(B|A)\mathbf{P}(A|B)\vec{P}(B). \tag{31}$$

Multiplying both sides by $\mathbf{P}(B|A)^{-1}$, we obtain

$$\mathbf{P}(B|A)^{-1}\vec{P}(B) = \mathbf{P}(A|B)\vec{P}(B). \tag{32}$$

Because the right-hand side is equal to $\vec{P}(A)$, we can multiply both sides by $\mathbf{P}(A)^{-1}$, obtaining

$$\mathbf{P}(A)^{-1}\mathbf{P}(B|A)^{-1}\vec{P}(B) = \vec{1}_A, \tag{33}$$

where $\vec{1}_A$ is the vector of 1s with support equal to the support of $\vec{P}(A)$. We define $\mathbf{R}$ by

$$\mathbf{P}(B|A)^{-1} = \mathbf{P}(A)\mathbf{R}\,\mathbf{P}(B)^{-1}. \tag{34}$$

Substituting equation (34) into (33), we obtain

$$\mathbf{R}\vec{1}_B = \vec{1}_A. \tag{35}$$

For any matrix $\mathbf{R}$ obeying (35) for all choices of $\mathbf{P}(B)$ and $\mathbf{P}(A|B)$, there is definition of $\mathbf{P}(B|A)$ that satisfies (18), given by

$$\mathbf{P}(B|A) = \mathbf{P}(B)\mathbf{R}^{-1}\mathbf{P}(A)^{-1}. \tag{36}$$

In retrospect, equation 36 is obvious. For all choices of $\mathbf{P}(A)$ and $\mathbf{P}(B)$, $\mathbf{P}(B|A)$ must map $\vec{P}(A)$ onto $\vec{P}(B)$. This requirement is clearly satisfied only if the matrix $\mathbf{R}$ given in (36) maps $\vec{1}_B$ onto $\vec{1}_A$. For the columns of $\mathbf{P}(B|A)$ to sum to 1, we further require[3]

$$\mathbf{P}(B|A)^T\vec{1}_B = \vec{1}_A. \tag{37}$$

[3]One example of a definition for $\mathbf{P}(B|A)$ that obeys (18) but has columns that do not sum to 1 can be obtained in the case of $N_A = N_B$ by setting $\mathbf{R}$ equal to $\mathbf{P}(A|B)^T$.

From (36), this is equivalent to the statement that

$$\vec{P}(A) = (\mathbf{R}^T)^{-1}\vec{P}(B),$$

or equivalently,

$$\mathbf{R}^T\vec{P}(A) = \vec{P}(B), \tag{38}$$

so $\mathbf{R}$ must itself be the transpose of a conditional probability distribution. The challenge is now to enumerate the set of matrices $\mathbf{R}$ that satisfy both (35) and (38). We can categorize the choices by the order of their dependency on $\mathbf{P}(A|B)$.

4.1. *Zeroth-order axioms.* It is possible to specify $\mathbf{P}(B|A)$ entirely independently of $\mathbf{P}(A|B)$, while still satisfying both of the requirements of (18). One such specification involves setting all of the columns of $\mathbf{P}(B|A)$ to be equal to $\vec{P}(B)$, or equivalently, setting $\mathbf{R}^{-1}$ to be the matrix in which all of the rows are given by $\vec{P}(A)$. However, this axiom is uninteresting from the perspective of inference, since it is impossible to infer $\vec{P}(B)$ without knowing $\vec{P}(B)$ to begin with. The corresponding rule says, if $\vec{P}(B)$ is already known, one can simply ignore all measurements.

4.2. *First-order axioms.* The set of first-order axioms is a one-dimensional convex hull, the two extreme points of which form a conjugate pair. The first extreme point is obtained by setting $\mathbf{R}$ equal to $(P(A|B)^T)^{-1}$. Applying (36), we recover Bayes' rule:

$$\mathbf{P}(B|A) = \mathbf{P}(B)\mathbf{P}(A|B)^T\mathbf{P}(A)^{-1}. \tag{39}$$

The matrix $\mathbf{P}(B|A)^T$ given in (39) also maps $\vec{1}_B$ onto $\vec{1}_A$, and also satisfies (38), so it is also a candidate for $\mathbf{R}$. This results in the second axiom:

$$\mathbf{R} = \mathbf{P}(A)^{-1}\mathbf{P}(A|B)\mathbf{P}(B) \tag{40}$$

Applying (36), we find the inversion rule,

$$\mathbf{P}(B|A) = \mathbf{P}(A|B)^{-1}. \tag{41}$$

The conjugacy of Bayes' rule and the inversion rule is now obvious. Denoting by $\mathbf{R}_K$ the axiom specifying Bayes' rule and by $\mathbf{R}_I$ the axiom specifying the inversion rule, we have

$$\mathbf{P}(B|A)_K^T = \mathbf{R}_I, \tag{42}$$

and likewise,

$$\mathbf{P}(B|A)_I^T = \mathbf{R}_K. \tag{43}$$

The set of all first-order axioms may then be obtained by noting that any convex combination of first-order axioms is also a first-order axiom, so the set of first-order axioms is a convex hull. Since there are no other matrix functions first order in $P(A|B)$ that obey (35) and (38), the entire set of first-order axioms may be parametrized by

$$\mathbf{R} = p\mathbf{P}(A)^{-1}\mathbf{P}(A|B)\mathbf{P}(B) + (1-p)(\mathbf{P}(A|B)^T)^{-1}, \tag{44}$$

for $p \in [0,1]$.

4.3. *Higher-order Axioms.* Just as any convex combination of $\mathbf{R}$ matrices is also a $\mathbf{R}$ matrix, any product of $\mathbf{R}$ matrices and $\mathbf{R}$-matrix inverses is a $\mathbf{R}$ matrix as well. To see this, note that if $\mathbf{R}$ maps $\vec{1}_B$ to $\vec{1}_A$, then $\mathbf{R}^{-1}$ maps $\vec{1}_A$ to $\vec{1}_B$. Likewise, if $\mathbf{R}^T$ maps $\vec{P}(A)$ to $\vec{P}(B)$, then $(\mathbf{R}^T)^{-1}$ maps $\vec{P}(B)$ to $\vec{P}(A)$. The set of higher-order axioms may thus be generated by combining different first-order $\mathbf{R}$ matrices in the pattern

$$\mathbf{R} = \mathbf{R}_1\mathbf{R}_2^{-1}\mathbf{R}_3\ldots, \tag{45}$$

corresponding to the rule

$$\mathbf{P}(B|A) = \mathbf{P}(B|A)_1\mathbf{P}(B|A)_2^{-1}\mathbf{P}(B|A)_3, \tag{46}$$

where $\mathbf{P}(B|A)_i$ is the conditional probability associated with $\mathbf{R}_i$. Evidently, rules only exist at odd orders, and clearly the set of nth-order rules is the convex hull of the nth-order product set of the set of first-order rules. As an example of a rule at higher-order, we have

$$\begin{aligned}\mathbf{R} &= \mathbf{R}_K\mathbf{R}_I^{-1}\mathbf{R}_K \\ &= (\mathbf{P}(A|B)^T)^{-1}\left(\mathbf{P}(B)^{-1}\mathbf{P}(A|B)^{-1}\mathbf{P}(A)\right)(\mathbf{P}(A|B)^T)^{-1},\end{aligned} \tag{47}$$

which implies

$$\mathbf{P}(B|A) = \mathbf{P}(B)\mathbf{P}(A|B)^T\mathbf{P}(A)^{-1}\mathbf{P}(A|B)\mathbf{P}(B)\mathbf{P}(A|B)^T\mathbf{P}(A)^{-1}. \tag{48}$$

This particular rule has the property that it discounts the evidence in favor of the prior and the model, as is evident from the fact that in this system, $\mathbf{P}(B|A)$ depends only at first order on $\mathbf{P}(A)$, while it depends on

$\mathbf{P}(B)$ at second order and $\mathbf{P}(A|B)$ at third order. Although it is more costly to implement computationally, it also displays some properties that may be useful for machine learning applications. For example, unlike for Bayes' rule or the inversion rule, the columns of $\mathbf{P}(B|A)$ will not in general sum to 1 unless the correct values of $\mathbf{P}(B)$ and $\mathbf{P}(A)$ are used in the calculation, thus providing a test of convergence.

5. Conclusion. We have shown that the axiomatization of probability theory in terms of set theory originally formulated by Kolmogorov imposes strong conditions on the form of the reverse conditional probabilities. These conditions ensure desirable properties required for consistency with the frequentist interpretation of probability theory, including positivity and correct behavior when $P(A, B)$ is a product distribution or not invertible. However, they are not necessary for correct statistical inference. Within a more general axiomatization, one can use alternative rules for statistical inference in parallel with Bayes' rule, leading to unexpected new ways of performing inference. As a result, alternative rules of inference may prove more useful than Bayes' rule for some applications.

In addition, we ask whether there may be physical circumstances in which an alternative axiomatization of probability theory would be preferred to the Kolmogorov theory. For example, it is known that although observable probabilities are always positive, the underlying distributions describing quantum mechanical states do not necessarily obey positivity Wigner [1932]. Indeed, the violation of positivity by quantum states has been shown to be a fundamental distinguishing property of quantum mechanics Spekkens [2008]. It may be that for describing these quantum mechanical probability distributions, which are not necessarily subject to a standard frequentist interpretation, an alternative axiomatization of probability is more suitable.

References.

Andrey N Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitstheorie*. Springer, Berlin, 1933.

Robert W Spekkens. Negativity and contextuality are equivalent notions of nonclassicality. *Physical review letters*, 101(2):020401, 2008.

Eugene Wigner. On the quantum correction for thermodynamic equilibrium. *Physical Review*, 40(5):749, 1932.

Synthetic Neurobiology Group
MIT Media Lab, 75 Amherst Street
Cambridge, MA, 02139
E-mail: sgr@mit.edu